arXiv:1609.06954v1 [cs.AI] 21 Sep 2016

Semiring Programming: A Framework for Search, Inference and Learning

Vaishak Belle
vaishak@cs.kuleuven.be

Luc De Raedt
luc.deraedt@cs.kuleuven.be

Abstract

To solve hard problems, AI relies on a variety of disciplines such as logic, probabilistic reasoning, machine learning and mathematical programming. Although it is widely accepted that solving real-world problems requires an integration amongst these, contemporary representation methodologies offer little support for this.

In an attempt to alleviate this situation, we introduce a new declarative programming framework that provides abstractions of well-known problems such as SAT, Bayesian inference, generative models, and convex optimization. The semantics of programs is defined in terms of first-order structures with semiring labels, which allows us to freely combine and integrate problems from different AI disciplines.

1 Introduction

AI applications, such as robotics and logistics, rely on a variety of disciplines such as logic, probabilistic reasoning, machine learning and mathematical programming. These applications are often described in a combination of natural and mathematical language, and need to be engineered for the individual application. Formal languages that allow declarative specifications enable re-use and descriptive clarity.

The subfields of AI have contributed many formalisms for the declarative modeling and effective solving of problems. In knowledge representation and constraint programming, for example, languages such ASP (Brewka, Eiter, and Truszczyński 2011) and Essence (Frisch et al. 2008) are prominent, which use SAT, SMT and MIP technology. In machine learning and probabilistic reasoning, statistical relational learning systems and probabilistic programming languages such as Markov Logic (Richardson and Domingos 2006), Church (Goodman et al. 2008) and Problog (De Raedt, Kimmig, and Toivonen 2007) are prominent, which use sampling and model counting methodologies. Finally, in mathematical programming and optimization, disciplined programming (Grant and Boyd 2008) and AMPL (Fourer, Gay, and Kernighan 1993) have been developed.

Be that as it may, it is widely accepted that solving real-world AI problems requires an integration of different disciplines. Consider, for example, that a robot may decide its course of action using a SAT-based planner, learn about the world using Kalman filters, and grasp objects using geometric optimization technology (Srivastava et al. 2014). But contemporary declarative frameworks offer little support for such universality: KR and constraint formalisms mostly focus on model generation for discrete problems, probabilistic programming languages do not handle linear and arithmetic constraints, and finally, optimization frameworks work with linear algebra and algebraic constraints to specify the problem and thus are quite different from the high-level descriptions used in the other disciplines and do not support probabilistic or logical reasoning.

What is lacking here is a universal modeling framework that allows us to declaratively specify problems involving *logic* and *constraints*, *mathematical programs*, as well as *discrete and continuous probability distributions* in a simple, uniform, modular and transparent manner. Such a framework would greatly simplify the development and understanding of AI systems with integrated capabilities, and would tame the model building process. In this paper, we propose a new declarative framework called *semiring programming* (SP) that is an attempt to fill this gap.

Eugene Freuder (1997) famously quipped: "constraint programming represents one of the closest approaches computer science has yet made to the Holy Grail of programming: the user states the problem, the computer solves it." The underlying idea was summarized in the slogan:

(constraint) program = model + solver

The vision of SP builds upon this equation in that:

(semiring) program = logical theory + semiring + solver

Thus a model is expressed in a logical system in tandem with a semiring and a weight function. In its essence, semiring programming rests on three important developments:

- Semirings have been successfully shown to unify specifications such as finding a solution versus counting the number of solutions (Bistarelli 2004).
- Weighted model counting represents perhaps the simplest approach to reason about weights on possible worlds and elegantly decouples the logical representation from the weight function (Chavira and Darwiche 2008).
- Probabilistic programming languages like Church can appeal to their host languages (e.g. LISP) to naturally support composition and abstraction.

As a framework, SP is set up to allow the modeler to freely choose the logical theory (syntax and semantics) and so everything from non-classical logical consequence to real-arithmetic is fair game. Together with a semiring and these weights, a program computes the *count*. Our task will be to show that the usual suspects from AI disciplines, such as SAT, CSP, Bayesian inference, and convex optimization can be: (a) expressed as a program, and (b) count is a solution to the problem, that is, the count can be a {*yes*, *no*} answer in SAT, a probability in Bayesian inference, or a bound in convex optimization. In other words, the count is shown to provide abstractions for challenging AI tasks. We will also demonstrate a variety of more complex problems, such as matrix factorization, and one on compositionality.

The fundamental question, then, is this: how are to we provide a semantics for such a programming language? It is not possible to piggyback on mathematical accounts of optimization problems in an obvious way, because the language is allowed to be richer than what is usually assumed. That is a central focus of this paper.

At the outset, we remark that our concern here is on the formal nature and generality of the framework, in terms of its modeling capabilities. Ideas on solvers are left to a discussion, and their realizations to the future.

2 Preliminaries

Logical Setup

The framework is developed in a general way, agnostic about the meaning of sentences. We adopt (and assume familiarity with) terminology from predicate logic (Enderton 1972).

Definition 1: A *theory* $\mathcal{T}$ is a triple $(\mathcal{L}, \mathcal{M}, \models)$, where $\mathcal{L}$ is a set of sentences called the *language* of the theory, $\mathcal{M}$ a set called the *models* of the theory, and $\models$ a subset of $\mathcal{M} \times \mathcal{L}$ called the *satisfiability relation*.

The set $\mathcal{L}$ is implicitly assumed to be defined over a vocabulary $vocab(\mathcal{L})$ of relation and function symbols, each with an associated arity. Constant symbols are 0-ary function symbols. Every $\mathcal{L}$-model $M \in \mathcal{M}$ is a tuple containing a universe $dom(M)$, and a relation (function) for each relation (function) symbol of $vocab(\mathcal{L})$. For relation (constant) symbol p, the relation (universe element) corresponding to p in a model M is denoted p^M. For $\phi \in \mathcal{L}$ and $M \in \mathcal{M}$, we write $M \models \phi$ to say that M *satisfies* ϕ. We let $\mathcal{M}(\phi)$ denote $\{M \in \mathcal{M} \mid M \models \phi\}$. Finally, the set $lits(\mathcal{L})$ denotes the literals in $\mathcal{L}$, and we write $l \in M$ to refer to the $\mathcal{L}$-literals that are satisfied at M.

This formulation is henceforth used to instantiate a particular logical system, such as fragments/extensions of first-order logic, as well as non-classical logical consequence.

Example 2: Suppose $\mathcal{T} = (\mathcal{L}, \mathcal{M}, \models)$ is as follows: $\mathcal{L}$ is obtained using a set of 0-ary predicates $\mathcal{P}$ and Boolean connectives, $\mathcal{M}$ are $\{0, 1\}$ assignments to the elements of $\mathcal{P}$, and $\models$ denotes satisfaction in propositional logic via the usual inductive rules. Then, we obtain a *propositional theory*.

Example 3: Define a theory $(\mathcal{L}, \mathcal{M}, \models)$ where $\mathcal{L}$ is the positive Horn fragment from a finite vocabulary, $\mathcal{M}$ the set of propositional interpretations, and for $M \in \mathcal{M}, \phi \in \mathcal{L}$, define $M \models \phi$ iff $M \models_{\text{PL}} \phi$ (*i.e.*, satisfaction in propositional logic) and for all $M' \subsetneq M, M' \not\models_{\text{PL}} \phi$. This theory can be used to reason about the *minimal models* of a formula.

Weighted Model Counting

Semiring programming draws from the conceptual simplicity of weighted model counting (WMC), which we briefly recap here. WMC is an extension of #SAT, where one simply counts the number of models of a propositional formula (Gomes, Sabharwal, and Selman 2009). In WMC, one accords a weight to a model in terms of the literals true at the model, and computes the sum of the weights of all models.

Definition 4: Suppose ϕ is a formula from a propositional language $\mathcal{L}$ with a finite vocabulary, and suppose $\mathcal{M}$ is the set of $\mathcal{L}$-models. Suppose w: $lits(\mathcal{L}) \to \mathbb{R}^{\geq 0}$ is a weight function. Then:[1]

$$\text{WMC}(\phi, w) = \sum_{M \models \phi} \prod_{l \in M} w(l)$$

is called the *weighted model count* (WMC) of ϕ.

The formulation elegantly decouples the logical sentence from the weight function. In this sense, it is clearly agnostic about how weights are specified in the modeling language, and thus, has emerged as an assembly language for Bayesian networks (Chavira and Darwiche 2008), and probabilistic programs (Fierens et al. 2011), among others.

Commutative Semirings

Our programming model is based on algebraic structures called *semirings* (Kuich 1997); the essentials are as follows:

Definition 5: A (commutative) semiring $\mathcal{S}$ is a structure $(\mathbb{S}, \oplus, \otimes, \mathbf{0}, \mathbf{1})$ where $\mathbb{S}$ is a set called the *elements* of the semiring, $\oplus$ and $\otimes$ are associative and commutative, $\mathbf{0}$ is the identity for $\oplus$, and $\mathbf{1}$ is the identity for $\otimes$.

Abusing notation, when the multiplication operator is not used, we simply refer to the triple $(\mathbb{S}, \oplus, \mathbf{0})$ as a semiring.

Example 6: The structure $(\mathbb{N}, +, \times, 0, 1)$ is a commutative semiring in that for every $a, b \in \mathbb{N}$, $a + 0 = a$, $a \times 1 = a$, $a + b = b + a$, and so on.

3 Semiring Programming

In essence, the semiring programming scheme is as follows:

- **Input**: a theory $\mathcal{T} = (\mathcal{L}, \mathcal{M}, \models)$, a sentence $\phi \in \mathcal{L}$, a commutative semiring $\mathcal{S}$, and a weight function w.
- **Output**: the *count*, denoted $\#(\phi, w)$.

The scope of these programs is broad, and so we will need different kinds of generality. Roughly, the distinction boils down to: (a) whether the set of models for a formula is *finite* or *infinite*, and (b) whether the weight function can be *factorized* over the literals or not (in which case the weight function directly labels the models of a theory). The thrust of this section is: (a) to show how these distinctions subsume important model generation notions in the literature,

[1] In the context of $\mathcal{L}$-models $\mathcal{M}$, we simply write $t = \sum_{M \models \phi} u$ to mean $t = \sum_{\{M \in \mathcal{M} \mid M \models \phi\}} u$.

	factorized	non-factorized
finite	WMC	logical-integer programming
infinite	polyhedron volume	convex optimization

Figure 1: scope of programs

```
(set-logic PL)
(set-algebra [NAT,max,*,0,1])
(declare-predicate p ())
(declare-predicate q ())
F = (p or q)
(declare-weight (p 1))
(declare-weight ((neg p) 2))
(declare-weight (q 3))
(declare-weight ((neg q) 4))
(count F)
```

Figure 2: most probable explanation

and (b) providing rigorous definitions for the count operator. In terms of organization, we begin with the finite case, before turning to the infinite ones. We present an early preview of some of the models considered in Figure 1.

Finite

Here, we generalize the WMC formulation to semiring labels, but also go beyond classical propositional logic.

Definition 7: We say the theory $\mathcal{T} = (\mathcal{L}, \mathcal{M}, \models)$ is *finite* if for every $\phi \in \mathcal{L}$, $\{M \in \mathcal{M} \mid M \models \phi\}$ is finite.

So, a propositional language with a finite vocabulary is a finite theory, regardless of (say) standard or minimal models. Similarly, a first-order language with a finite Herbrand base is also a finite theory.

Definition 8: Suppose $\mathcal{T} = (\mathcal{L}, \mathcal{M}, \models)$ is a finite theory. Suppose $\mathcal{S} = (\mathbb{S}, \oplus, \otimes, \mathbf{0}, \mathbf{1})$ is a commutative semiring. Suppose $w : \mathcal{M} \to \mathbb{S}$. For any $\phi \in \mathcal{L}$, we define:

$$\#(\phi, w) = \bigoplus_{M \models \phi} w(M)$$

If $w : \mathit{lits}(\mathcal{L}) \to \mathbb{S}$, then the problem is *factorized*, where:

$$w(M) = \bigotimes_{l \in M} w(l).$$

Essentially, as in WMC, we sum over models and take products of the weights on literals but wrt a particular semiring.

Example 9: We demonstrate SAT and #SAT. Consider a propositional theory $\mathcal{T} = (\mathcal{L}, \mathcal{M}, \models)$ where $\mathit{vocab}(\mathcal{L}) = \{p, q\}$, and suppose $\phi = (p \vee q)$. Letting $\mathcal{M}$ be standard $\mathcal{L}$-models, clearly $|\mathcal{M}| = 4$ and $|\{M \in \mathcal{M} \mid M \models \phi\}| = 3$. Suppose for every $M \in \mathcal{M}$, w is function such that $w(M) = 1$. For the semiring $(\{0, 1\}, \vee, 0)$:

$$\#(\phi, w) = \bigvee_{M \models \phi} w(M) = 1 \vee 1 \vee 1 = 1.$$

Consider the semiring $(\mathbb{N}, +, 0)$ instead. Then

$$\#(\phi, w) = \sum_{M \models \phi} w(M) = 1 + 1 + 1 = 3.$$

As with Definition 4, the framework is agnostic about the modeling language. But for presentation purposes, programs are sometimes described using a notation inspired by the SMT-LIB standard (Barrett, Fontaine, and Tinelli 2015).

Example 10: We demonstrate MPE (most probable explanation) and WMC. Consider the theory, semiring and weight function w from Figure 2, which specifies, for example, a vocabulary of two propositions p and q, $w(p) = 1$ and $w(\neg p) = 2$. In accordance with that semiring, the weight of a model, say $\{p, \neg q\}$, of the formula `F` is $1 \times 4 = 4$. Thus, for the semiring $(\mathbb{N}, \max, \times, 0, 1)$, we have:

$$\#(\mathtt{F}, w) = \max\{6, 3, 4\} = 6$$

which finds the most probable assignment. Consider the semiring $(\mathbb{N}, +, \times, 0, 1)$ instead. Then:

$$\#(\mathtt{F}, w) = 6 + 3 + 4 = 13$$

which gives us the weighted model count.

Example 11 : Extending the discussion in (Hooker and Lama 1999), we consider a class of mathematical programs where linear constraints and propositional formulas can be combined freely. See Figure 3 for an example with non-linear objectives. Formally, quantifier-free linear integer arithmetic and propositional logic are specified as the underlying logical systems, and the domains of constants are typed. The program declares formulas `F`, `G`, and `H`.

The counting task is non-factorized, and our convention for assigning weights to models is by letting the `declare-weight` directive also take arbitrary formulas as arguments. Of course, `TRUE` holds in every model, and so, the weight of every model is determined by the evaluation of `x1 * x2` at the model, that is, for any M, its weight is $*^M(\mathtt{x1}^M, \mathtt{x2}^M)$. For example, a model that assigns `1` to `x1` and `1` to `x2` is accorded the weight $1*1$. Computing the count over $(\mathbb{N}, \max, 0)$ then yields a model of `H` with the highest value for `x1 * x2`.

To see this program in action, consider that every model of `H` must satisfy `p2`, and so must admit `3 * x1 <= 4` and `2 * x2 <= 5`. Since $\{\mathtt{x1}, \mathtt{x2}\}$ can only take values $\{\mathtt{1}, \ldots, \mathtt{10}\}$, given the constraints, the desired model must assign `x1` and `x2` to `1` and `2` respectively. Then, its weight is $1 * 2$.

Encoding finite domain constraint satisfaction problems as propositional satisfiability is well-known. The benefit, then, of appealing to our framework is the ability to easily formulate counting instances:

Theorem 12: *Suppose Q is a CSP over variables $\mathcal{X}$, domains $\mathcal{D}$ and constraints C. There is a first-order theory $(\mathcal{L}, \mathcal{M}, \models)$, a $\mathcal{L}$-sentence ϕ and a weight function w such that $\#(\phi, w) = 1$ over $(\{0, 1\}, \vee, 0)$ iff Q has a solution. Furthermore, Q has n solutions iff $\#(\phi, w) = n$ over $(\mathbb{N}, +, 0)$.*

```
(set-logic QF_LIA;PL)
(set-algebra [NAT,max,0])
(set-type INT={1,...,10})
(declare-function x1 () INT)
(declare-predicate p1 ())
... /* declare x2 and p2 analogously */
F = ((p1 or p2) => 3*x1 <= 4)
G = (p2 => (2*x2 <= 5))
H = ((F and G) and p2)
(declare-weight TRUE x1*x2)
(count H)
```

Figure 3: logical-integer programming

Constraints are Boolean-valued functions (Freuder and Mackworth 2006), and so constraints over $\mathcal{X}$ can be encoded as $\mathcal{L}$-sentences. We omit the details, but do an example.

Example 13: See Figure 4 for a counting instance of graph coloring: `edge(x,y)` determines there is an edge between `x` and `y`, `node(x)` says that `x` is a node, and `color(x,y)` says that node `x` is assigned the color `y`. The actual graph is provided using the formula `DATA`, which declares a fully connected 3-node graph. Also, `CONS` is a conjunction of the usual coloring constraints, *e.g.*, an edge between nodes `x` and `y` means that they cannot be assigned the same color.

Let $\mathcal{M}$ be a set of first-order structures for the vocabulary {`edge`, `node`, `color`}, respecting types from Figure 4. The interpretation of {`edge`, `node`} is assumed to be the same for all the models in $\mathcal{M}$ and is as given by `DATA`. Basically, then, the models differ in their interpretation of `color`. One model of ϕ = (`DATA and CONS`), for example, is {`color(1,r)`, `color(2,g)`, `color(3,b)`}. The weights of all models is 1, and so, for $(\mathbb{N}, +, 0)$ we get:[2]

$$\#(\phi, w) = 6.$$

To summarize, the following result is easily shown for semiring programs:[3]

Theorem 14: *Suppose $\theta \in$ {SAT, #SAT, WMC, CSP, #CSP}. Suppose θ° is a solution to θ, in that $\theta^\circ \in \{0, 1\}$ for* SAT *and* CSP*, and $\theta^\circ \in \mathbb{R}$ for the rest. Then for any θ, there is a $\mathcal{T} = (\mathcal{L}, \mathcal{M}, \models)$, $\mathcal{S}$, w and $\phi \in \mathcal{L}$ such that $\#(\phi, w) = \theta^\circ$.*

Infinite: Non-factorized

Defining measures (Halmos 1950) on the predicate calculus is central to logical characterizations of probability theory (Halpern 1990; Chistikov, Dimitrova, and Majumdar 2015). We adapt this notion for semirings to introduce a general form of counting. For technical reasons, we assume that the universe of the semirings is $\mathbb{R}$. (If required, the range of the weight function can always be restricted to any subset of $\mathbb{R}$.)

[2]In an analogous fashion, soft and weighted CSPs can be expressed using semirings (Bistarelli et al. 1999).

[3]We remark that although the count operator in itself does not provide the variable assignments for SAT and CSP, we assume this can be retrieved from the satisfying interpretations.

```
(set-logic FOL)
(set-algebra [NAT,+,0])
(set-type COLOR={r,b,g})
(set-type NODE={1,2,3})
(declare-predicate node (NODE))
(declare-predicate edge (NODE,NODE))
(declare-predicate color (NODE,COLOR))
N = (node(1) and node(2) and node(3))
E = (edge(1,2) and edge(2,3) and edge(3,1))
DATA = (N and E)
CONS = /* coloring constraints (omitted) */
(declare-weight TRUE 1)
(count (DATA and CONS))
```

Figure 4: counting graph coloring instances

Definition 15: Let $\mathcal{S} = (\mathbb{R}, \oplus, \mathbf{0})$ be any semiring and $\mathcal{T} = (\mathcal{L}, \mathcal{M}, \models)$ a theory. Let Σ be a σ-algebra over $\mathcal{M}$ in that $(\mathcal{M}, \Sigma)$ is a σ-finite measurable space wrt the measure $\mu\colon \Sigma \to \mathbb{R}$ respecting $\oplus$. That is, for all $E \in \Sigma$, $\mu(E) \geq 0$, $\mu(\{\}) = 0$ and μ is closed under complement and countable unions: for all pairwise disjoint countable sets $E_1, \ldots \in \Sigma$, we have

$$\mu(\bigcup_{j=1}^{\infty} E_j) = \bigoplus_{j=1}^{\infty} \mu(E_j).$$

Moreover, because the spaces are σ-finite, $\mathcal{M}$ is the countable union of measurable sets with finite measure. Then for any $\phi \in \mathcal{L}$, $\#(\phi, \mu) = \mu(\mathcal{M}(\phi))$.

Example 16: We demonstrate convex optimization. Suppose $\mathcal{T} = (\mathcal{L}, \mathcal{M}, \models)$ is a first-order theory only containing constants $\mathcal{X} = \{x_1, \ldots, x_k\}$ with domains $D_i = \mathbb{R}$. Suppose $\phi(x_1, \ldots, x_k) \in \mathcal{L}$ is a conjunction of formulas of the form $c_1 x_1 + \ldots + c_k x_k \leq d$, for real numbers $c_1, \ldots, c_k, d$, and thus describing a polyhedron. In other words, for every $M \in \mathcal{M}$, $dom(M) = \mathbb{R}$, and so, M is a real-valued assignment to $\mathcal{X}$. In particular, if $M \models \phi$, then $x_1^M, \ldots, x_k^M$ is a point inside the polyhedron ϕ. Let Σ be a σ-algebra over $\mathcal{M}$ and so every $E \in \Sigma$ is a measurable set of points.

Suppose $f(x_1, \ldots, x_k)\colon \mathbb{R}^k \to \mathbb{R}$ is a convex function that we are to minimize. Consider the semiring $(\mathbb{R}, \inf, 0)$ and a measure μ such that for any $E \in \Sigma$:

$$\mu(E) = \inf \left\{ f(x_1^M, \ldots, x_k^M) \mid M \in E \right\}$$

which finds the infimum of the f-values across the assignments in E. Then, $\mu(\mathcal{M}(\phi)) = \#(\phi, \mu)$ gives the minimum of the convex function in the feasible region determined by ϕ.

Suppose f is a concave function that is to be maximized. We would then use $(\mathbb{R}, \sup, 0)$ instead, which finds the supremum of the f-values across assignments in $E \in \Sigma$.

More generally, the same construction is easily shown to be applicable for other families of mathematical programming (*e.g.*, non-linear optimization) as follows:

Theorem 17: *Suppose $\{x_1, \ldots, x_k\}$ is a set of real-valued variables. Suppose $P \subset \mathbb{R}^k$ is the feasible region of an optimization problem of the form $g_i(x_1, \ldots, x_k) \circ d_i$ for $i \in \mathbb{N}$,*

```
(set-logic LRA)
(set-algebra [REAL,inf,0])
(declare-function input (INT,INT) REAL)
... /* declare left, right, app */
(declare-function err () REAL)
DATA = /* entries in input matrix (omitted) */
F = app(x,y) == sum{e} left(x,e)*right(e,y)
G = err == norm(sum{x,y} input(x,y) - app(x,y))
(declare-weight TRUE err)
(count (DATA and F and G))
```

Figure 5: matrix factorization

where $\circ \in \{\leq, <, \geq, >\}$ *and* $g_i : \mathbb{R}^k \to \mathbb{R}$. *Suppose* $f : \mathbb{R}^k \to \mathbb{R}$ *is a function to be maximized (minimized). Then there is a* $\mathcal{T} = (\mathcal{L}, \mathcal{M}, \models), \mathcal{S}, \mu$ *and* $\phi \in \mathcal{L}$ *such that* $\#(\phi, \mu)$ *is the maximum (minimum) value for* f *in the feasible region P.*

Example 18 : For a non-trivial example, consider the problem of matrix factorization, a fundamental concern in information retrieval and computer vision (Ding, Li, and Jordan 2010). Given a matrix $I \in \mathbb{R}^{p \times n}$, we are to compute matrices $L \in \mathbb{R}^{p \times k}$ and $R \in \mathbb{R}^{n \times k}$, such that $e = \|I - LR^T\|$ is minimized. (Here, $\|\cdot\|$ denotes the Frobenius Norm.) Using real arithmetic, we provide a formulation in Figure 5. (Free variables are assumed to be implicitly quantified from the outside.)

Let $\mathcal{T} = (\mathcal{L}, \mathcal{M}, \models)$ be the theory of real arithmetic, where $\mathcal{L}$ includes the following function symbols: {`input`, `left`, `right`, `app`, `err`}. Here, `input(x,y)` is a real-valued function such that `x` $\in \{1, \ldots, p\}$ and `y` $\in \{1, \ldots, n\}$ in that `input(m,n)` is the entry at the m^{th} row and the n^{th} column of the matrix I; these entries are specified by `DATA`. Letting ϕ = (`DATA and F and G`), the set $\{M \in \mathcal{M} \mid M \models \phi\}$ are those $\mathcal{L}$-models whose interpretation of `input` is fixed by `DATA`. Basically, these models vary in their interpretations of {`left`, `right`}, which determines their interpretations for `app` and `err`. Here, `app` computes the product of the matrices `left` and `right`, and `err` computes the Frobenius Norm wrt `app` and `input`.

Let Σ be a σ-algebra over $\mathcal{M}$. The weight function in Figure 5 determines a measure μ such that for any $E \in \Sigma$:

$$\mu(E) = \inf \left\{ \texttt{err}^M \mid M \in E \right\}.$$

Therefore, $\#(\phi, \mu)$ yields the lowest `err` value; the model M such that $\texttt{err}^M = \#(\phi, \mu)$ is one with the best factorization of matrix I.

Infinite: Factorized

Despite the generality of the above definition, we would like to address the factorized setting for a number of applications, the most prominent being probabilistic inference in hybrid graphical models (Belle, Passerini, and Van den Broeck 2015). Consider, for example, a joint distribution on the probability space $\mathbb{R} \times \{0, 1\}$. Here, it is natural to define weights for each random variable separately, prompting a factorized formulation of counting. More generally, in many robotic applications, such hybrid spaces are common (Thrun, Burgard, and Fox 2005). The main idea then is to apply our definition for counting by measures to each variable independently, and construct a measure for the entire space by *product measures* (Halmos 1950).

A second technicality is that in the finite case, the set of literals true at a model was finite by definition. This is no longer the case. For example, suppose x is a real-valued variable in a language $\mathcal{L}$, and M is a $\mathcal{L}$-model that assigns 3 to x. Then, $M \models (x = 3)$ but also $M \models (x \neq 3.1)$, $M \models (x \neq 3.11), \ldots$, and so on. Thus, for technical reasons, we assume that $\mathcal{L}$ only consists of constant symbols $\mathcal{X} = \{x_1, \ldots, x_k\}$ with fixed (possibly infinite) domains $\{D_1, \ldots, D_k\}$; the measures are defined for these domains.

Definition 19 : Let $\mathcal{S} = (\mathbb{R}, \oplus, \otimes, \mathbf{0}, \mathbf{1})$ be any semiring, $\mathcal{T} = (\mathcal{L}, \mathcal{M}, \models)$ any theory where $vocab(\mathcal{L}) = \{x_1, \ldots, x_k\}$ over fixed domains D_i. Suppose $\phi \in \mathcal{L}$. Suppose Σ_i is a σ-algebra over D_i in that (D_i, Σ_i) is a σ-finite measurable space wrt the measure $\mu_i : \Sigma_i \to \mathbb{R}$ respecting $\oplus$ (as in Definition 15). Define the product measure $\mu^* \doteq \mu_1 \times \cdots \times \mu_k$ on the measurable space $(D_1 \times \cdots \times D_k, \Sigma_1 \times \cdots \Sigma_k)$ satisfying:[4]

$$\mu^*(E_1 \times \cdots \times E_k) = \mu_1(E_1) \otimes \cdots \otimes \mu_k(E_k)$$

for all $E_i \in \Sigma_i$. Finally, define

$$\#(\phi, \mu^*) = \mu^*([\phi])$$

where $[\phi] = \left\{ x_1^M \times \ldots \times x_k^M \mid M \in \mathcal{M}(\phi) \right\}$.

Intuitively, $E_1 \in \Sigma_1, \ldots, E_k \in \Sigma_k$ capture sets of assignments, and the product measure considers the algebraic product of the weights on assignments to terms. As before, for $E_i, E_i' \in \Sigma_i$, $\mu_i(E_i \cup E_i') = \mu_i(E_i) \oplus \mu_i(E_i')$. Finally, precisely because the measures are defined on the domains of the terms, we obtain these for all of the satisfying interpretations using the construction $[\phi]$.

Example 20 : We demonstrate the problem of finding the volume of a polyhedron, needed in the static analysis of probabilistic programs (Sankaranarayanan, Chakarov, and Gulwani 2013; Chistikov, Dimitrova, and Majumdar 2015). Suppose $\mathcal{T} = (\mathcal{L}, \mathcal{M}, \models)$ and $\phi \in \mathcal{L}$ is as in Example 16, that is, ϕ defines a polyhedron. For every 0-ary function symbol $x_i \in \{x_1, \ldots, x_k\}$ with domain $D_i = \mathbb{R}$, let Σ_i be the set of all Borel subsets of $\mathbb{R}$, and let μ_i be the Lebesgue measure. Thus, for any $E \in \Sigma_i$, $\mu_i(E)$ gives the length of this line. Then, for the semiring $(\mathbb{R}, +, \times, 0, 1)$, the $+$ operator sums the lengths of lines for each variable, and $\times$ computes the products of these lengths. Thus, $\#(\phi, \mu^*)$ is the volume of ϕ.

To see this in action, suppose $\phi = (2x \leq 5) \wedge (x \geq 1) \wedge (0 \leq y \leq 2)$. Then $\mathcal{M}(\phi) = \left\{ \langle x \to n, y \to m \rangle \mid n, m \in \mathbb{R}, \models \phi_{n,m}^{x,y} \right\}$, that is, all assignments to x and y such that $\phi_{n,m}^{x,y}$ is a valid expression in arithmetic. Therefore,

$$[\phi] = \{(n, m) \mid n \in [1, 2.5], m \in [0, 2], n \in \mathbb{R}, m \in \mathbb{R}\}.$$

[4]The product measure is unique owing to the σ-finite assumption via the Hahn-Kolmogorov theorem (Halmos 1950).

Assuming μ_x is the Lebesgue measure for all Borel subsets of the domain of x, we have

$$\mu_x(\{n \mid n \in [1, 2.5], n \in \mathbb{R}\}) = 1.5.$$

Analogously, $\mu_y(\{m \mid m \in [0, 2], m \in \mathbb{R}\}) = 2$. Then, the volume is $\#(\phi, \mu^*) = \mu^*([\phi]) = 1.5 \times 2 = 3$.

Example 21: We demonstrate probabilistic inference in hybrid models (Belle, Passerini, and Van den Broeck 2015) by extending Example 20. Consider a probabilistic program:

$X \sim$ UNIFORM$(0, 1)$
$Y \sim$ FLIP$(0, 1)$
if $(X > .6$ AND $Y \neq 1)$ **then** return DONE

In English: X is drawn uniformly from [0,1] and $Y \in \{0, 1\}$ is the outcome of a coin toss. If $X > .6$ and Y is not 1, the program terminates successfully. Suppose we are interested in the probability of DONE, which is expressed as the formula:

$$\phi = (0 \leq X \leq 1) \wedge (Y = 0 \vee Y = 1) \wedge (X > .6 \wedge Y \neq 1).$$

Suppose $\mathcal{T} = (\mathcal{L}, \mathcal{M}, \models)$ is the theory of linear real arithmetic, with $vocab(\mathcal{L}) = \{X, Y\}$, $D_X = \mathbb{R}$ and $D_Y = \{0, 1\}$. As in Example 20, let Σ_X be the set of all Borel subsets of $\mathbb{R}$ and μ_X the Lebesgue measure. Let Σ_Y be the set of all subsets of $\{0, 1\}$ and μ_Y be the counting measure, *e.g.,* $\mu_Y(\{\}) = 0, \mu_Y(\{1\}) = 1$, and $\mu_Y(\{0, 1\}) = 2$. So $\mathcal{M}(\phi) = \left\{\langle X \to n, Y \to m\rangle \mid n \in \mathbb{R}, m \in \{0, 1\}, \models \phi^{X,Y}_{n,m}\right\}$ and:

$$[\phi] = \{(n, m) \mid n > .6, 0 \leq n \leq 1, n \in \mathbb{R}, m \in \{0, 1\}\} \cap \{(n, m) \mid 0 \leq n \leq 1, n \in \mathbb{R}, m \neq 1, m \in \{0, 1\}\}.$$

This means that $\mu_X(\{n \mid \exists m \ (n, m) \in [\phi]\}) = .4$ and also $\mu_Y(\{m \mid \exists n \ (n, m) \in \phi\}) = 1$. Thus, $\mu^*([\phi]) = .4 \times 1$. Analogously, $\mu^*([\phi \vee \neg\phi]) = 2$. Therefore, the probability of DONE is $.4/2 = .2$.

In general, we have a variant of Theorem 17 (Hindi 2004):

Theorem 22: *Suppose $\{x_1, \ldots, x_k\}$ is any set of real-valued variables. Suppose D is any countable set. Suppose $P \subset \mathbb{R}^m \times D^n$, where $m + n = k$, is any region given by conjunctions of expressions of the form $c_1x_1 + \ldots + c_kx_k \circ e$ and $x_i \diamond d$, where $\circ \in \{\leq, <, \geq, >\}$, $\diamond \in \{\neq, =\}$, $c_i, e \in \mathbb{R}$ and $d \in D$. Then there is a $\mathcal{T} = (\mathcal{L}, \mathcal{M}, \models), \mathcal{S}, \mu$ and $\phi \in \mathcal{L}$ such that $\#(\phi, \mu)$ is the volume of P.*

4 Compositionality

A noteworthy feature of many logic-based knowledge representation formalisms is their compositional nature. In semiring programming, using the expressiveness of predicate logic, it is fairly straightforward to combine theories over possibly different signatures (*e.g.,* propositional logic and linear arithmetic), as seen, for example, in SMT solvers (Barrett et al. 2009).

A more intricate flavor of compositionality is when the new specification becomes difficult (or impossible) to define using the original components. This is a common occurrence in large software repositories, and has received a lot of attention in the AI community (Lierler and Truszczynski 2015).

In this section, we do not attempt to duplicate such efforts, but propose a different account of compositionality that is closer in spirit to semiring programming. It builds on similar ideas for CSPs (Bistarelli 2004), and is motivated by machine learning problems where learning (*i.e.,* optimization) and inference (*i.e.,* model counting) need to be addressed in tandem. More generally, the contribution here allows us to combine two semiring programs, possibly involving different semirings. For simplicity of presentation, we consider non-factorized and finite problems.

Definition 23 : Suppose $\mathcal{S}_1 = (\mathbb{S}_1, \oplus_1, \mathbf{0}_1)$ and $\mathcal{S}_2 = (\mathbb{S}_2, \oplus_2, \mathbf{0}_2)$ are any two semirings. We define their *composition* $\mathcal{S} = (\mathbb{S}, \oplus, \mathbf{0})$ as:

- $\mathbb{S} = \{(a, b) \mid a \in \mathbb{S}_1, b \in \mathbb{S}_2\}$;
- $\mathbf{0} = (\mathbf{0}_1, \mathbf{0}_2)$;
- for every $c = (a, b) \in \mathbb{S}$ and $c' = (a', b') \in \mathbb{S}$, let $c \oplus c' = (a \oplus_1 a', b \oplus_2 b')$.

That is, the composition of the semirings is formed from the Cartesian product, respecting the summation operator for the individual structures.

Definition 24 : Suppose $\mathcal{T}_1 = (\mathcal{L}_1, \mathcal{M}_1, \models_1)$ and $\mathcal{T}_2 = (\mathcal{L}_2, \mathcal{M}_2, \models_2)$ are theories, where $\mathcal{L}_1$ and $\mathcal{L}_2$ do not share atoms, $\mathcal{S}_1$ and $\mathcal{S}_2$ semirings, and w_1 and w_2 weight functions for $\mathcal{M}_1$ and $\mathcal{M}_2$ respectively. Given the environments $(\mathcal{T}_1, \mathcal{S}_1, w_1)$ and $(\mathcal{T}_2, \mathcal{S}_2, w_2)$, we define its *composition* as $(\mathcal{T}, \mathcal{S}, w)$, where $\mathcal{S}$ is a composition of $\mathcal{S}_1$ and $\mathcal{S}_2$ and $\mathcal{T} = (\mathcal{L}, \mathcal{M}, \models)$:

- $\phi \in \mathcal{L}$ is obtained over Boolean connectives from $\mathcal{L}_1 \cup \mathcal{L}_2$.
- $\mathcal{M} = \{(M_1, M_2) \mid M_i \in \mathcal{M}_i\}$.
- The meaning of $\phi \in \mathcal{L}$ is defined inductively:
  - $(M_1, M_2) \models p$ for atom p iff $M_1 \models_1 p$ if $p \in \mathcal{L}_1$ and $M_2 \models_2 p$ otherwise;
  - $(M_1, M_2) \models \neg\phi$ iff $(M_1, M_2) \not\models \phi$;
  - $(M_1, M_2) \models \phi_x \wedge \phi_y$ iff $(M_1, M_2) \models \phi_i$ for $i \in \{x, y\}$.
- For any $(M_1, M_2) \in \mathcal{M}$, $w((M_1, M_2)) = (w_1(M_1), w_2(M_2))$.

In essence, the Cartesian product for the semirings is extended for arbitrary theories and weight functions. The meaning of formulas rests on the property that $\mathcal{L}_1$ and $\mathcal{L}_2$ do not share atoms.[5] It is now easy to see that the counting for problem for $\mathcal{T}$ works as usual: that is, for any $\phi \in \mathcal{L}$,

$$\#(\phi, w) = \bigoplus_{(M_1, M_2) \models \phi} w((M_1, M_2)).$$

Example 25: We demonstrate a (simple) instance of combined learning and inference. Imagine a robot navigating a world by performing *move* actions, and believes its actuators need repairs. But before it alerts the technician, it would like to test this belief. A reasonable test, then, is to inspect its trajectory so far, and check whether the expected outcome of a move action in the current state matches the behavior of

[5] This is analogous to SMT solvers for combinations of theories (Barrett et al. 2009). However, see (Lierler and Truszczynski 2015; Ensan and Ternovska 2015) for accounts of modularity based on first-order structures sharing vocabularies (and thus, atoms).

the very first move action. More precisely, the robot needs to appeal to linear regression to estimate its expected outcome, and query its beliefs based on the regression model. We proceed as follows.

Environment 1 Let $\mathcal{T}_1 = (\mathcal{L}_1, \mathcal{M}_1, \models_1)$ be the theory of real arithmetic with $vocab(\mathcal{L}_1) = \{s_0, s_1, \ldots, s_k, a, b, e\}$. Suppose that by performing a *move* action, the robot's position changes from s_i to s_{i+1}. Let $\phi_1 \in \mathcal{L}_1$ be as follows:

$$(s_0 = 1 \wedge s_1 = 2 \wedge s_2 = 3) \wedge$$
$$e = \textstyle\sum_i (s_{i+1} - b - a \cdot s_i)^2$$

The idea is that the values of s_i are the explanatory variables in the regression model and s_{i+1} are the response variables, that is, the trajectory data is of the form $\{(s_0, s_1), (s_1, s_2)\}$.

Consider a weight function $w_1 : M \rightarrow e^M$. That is, the weight of a model M is the universe element corresponding to the constant e in M, analogous to Figure 3. (For simplicity, we assume that the coefficients of the regression model are natural numbers.) For the semiring $\mathcal{S}_1 = (\mathbb{N}, \min, \infty)$:

$$\#(\phi_1, w_1) = \min \left\{ e^M \mid M \models \phi_1, M \in \mathcal{M}_1 \right\}.$$

In other words, models in $\mathcal{M}_1$ interpret s_i as given by the data, and models differ in their interpretation of a, b and thus, e. For the data in ϕ_1, we would have a model M where $e^M = 0$, $a^M = 1$ and $b^M = 1$, and so $\#(\phi_1, w_1) = 0$.

Environment 2 Let $\mathcal{T}_2 = (\mathcal{L}_2, \mathcal{M}_2, \models_2)$ be a propositional theory, where $vocab(\mathcal{L}_2) = \{repair\}$. Imagine a weight function w_2 as follows: $w_2(\{repair\}) = .7$, and $w_2(\{\neg repair\}) = .3$. That is, the robot believes that repairs are needed with a higher probability. Indeed, for the semiring $(\mathbb{R}, +, 0)$:

$$\#(repair, w_2) = .7 \text{ versus } \#(\neg repair, w_2) = .3.$$

Composition Let $(\mathcal{T}, \mathcal{S}, w)$ be the composition of the two environments with $\mathcal{T} = (\mathcal{L}, \mathcal{M}, \models)$ and $\mathcal{S} = (\mathbb{S}, \oplus, \mathbf{0})$. Suppose $\phi \in \mathcal{L}$ is as follows:

$$\phi_1 \wedge (s_3 = a \cdot s_2 + b) \wedge$$
$$(s_3 - s_2) \neq (s_1 - s_0) \equiv repair.$$

The final conjunct basically checks whether the expected change in position matches what was happening initially, and if not, *repair* should be true.

To see Definition 24 in action, observe that for any $M_1 \in \mathcal{M}_1, M_2 \in \mathcal{M}_2$, $(M_1, M_2) \in \mathcal{M}$, the formulas ϕ_1 and those involving s_i are interpreted in M_1, but *repair* in M_2. The weight function is as follows. Given $(M_1, M_2) \in \mathcal{M}$ and $(M_1', M_2') \in \mathcal{M}$, we have:

$$w((M_1, M_2)) \oplus w((M_1', M_2'))$$
$$= (\ \min(w_1(M_1), w_1(M_1')),\ w_2(M_2) + w_2(M_2')\ ).$$

Then the robot can obtain the weight of *repair* and ϕ using:

$$\#(\phi \wedge repair, w) = (0, 0)$$

where, of course, the first argument is the error of the regression model and the second is 0 because $\phi \wedge repair$ is inconsistent. That can be contrasted to the count below:

$$\#(\phi \wedge \neg repair, w) = (0, .3).$$

It is also easy to see that $\#(\phi, w) = (0, .3)$.

As in WMC (Chavira and Darwiche 2008), suppose the robot obtains the probability of a query q given ϕ using:

$$\frac{\#(\phi \wedge q, w)}{\#(\phi, w)}$$

where the division is carried out by ignoring the regression error. Then the probability of $\neg repair$ given ϕ is 1. Thus, no repairs are needed.

5 Solver Construction

The upshot of semiring programming is that it encourages us to inspect strategies for a unified inferential mechanism (Kautz 2004). This has to be done carefully, as we would like to build on scalable methodologies in the literature, by restricting logical theories where necessary. In this section, we discuss whether our programming model can be made to work well in practice.

Let us consider two extremes:

- **Option 1:** At one extreme is a solver strategy based on a single computational technique. Probabilistic programming languages, such as Church (Goodman et al. 2008), have made significant progress in that respect for generative stochastic processes by appealing to Markov Chain Monte Carlo sampling techniques. Unfortunately, such sampling techniques do not scale well on large problems and have little support for linear and logical constraints.
- **Option 2:** At the other extreme is a solver strategy that is arbitrarily heterogeneous, where we develop unique solvers for specific environments, that is, $(\mathcal{T}, \mathcal{S})$ pairs.

Option 3: We believe the most interesting option is in between these two extremes. In other words, to identify the smallest set of computational techniques, and effectively integrate them is both challenging and insightful. This may mean that such a strategy is less optimal than Option 2 for the environment, but we would obtain a simpler and more compact execution model. To that end, let us make the following observations from our inventory of examples:

- **Finite versus infinite:** variable assignments are taken from finite sets versus infinite or uncountable sets.
- **Non-factorized versus factorized:** the former is usually an optimization problem with an objective function that is to be maximized or minimized. The latter is usually a counting problem, where we would need to identify one or all solutions.
- **Compositionality:** locally consistent solutions (*i.e.*, in each environment $(\mathcal{T}, \mathcal{S})$) need to be tested iteratively for global consistency.

Thus, Option 3 would be realized as follows:

- Factorized problems need a methodology for effective enumeration, and therefore, advances in model counting (Gomes, Sabharwal, and Selman 2009), such as knowledge compilation, are the most relevant. For finite theories, we take our cue from the Problog family of languages (Fierens et al. 2011; Kimmig, Van den Broeck, and De Raedt 2011), that

have effectively applied arithmetic circuits for tasks such as WMC and MPE. In particular, it is shown in (Kimmig, Van den Broeck, and De Raedt 2012) how arbitrary semiring labels can be propagated in the circuit. See (Fargier, Marquis, and Niveau 2013; Fargier, Marquis, and Schmidt 2013) for progress on knowledge compilation in CSP-like environments.

For infinite theories, there is growing interest in effective model counting for linear arithmetic using SMT technology (Chistikov, Dimitrova, and Majumdar 2015; Belle, Passerini, and Van den Broeck 2015; Belle, Passerini, and Van den Broeck 2016). Like in (Kimmig, Van den Broeck, and De Raedt 2011), however, we would need to extend these counting approaches to arbitrary semirings.

- For non-factorized problems, a natural candidate for handling semirings does not immediately present itself, making this is a worthwhile research direction.[6] Appealing to off-the-shelf optimization software (Andrei 2013) is always an option, but they embody diverse techniques and the absence of a simple high-level solver strategy makes adapting them for our purposes less obvious. In that regard, solvers for *optimization modulo theories* (OMT) (Sebastiani and Tomasi 2015) are perhaps the most promising. OMT technology extends SMT technology in additionally including a cost function that is be maximized (minimized). In terms of expressiveness, CSPs (Nieuwenhuis and Oliveras 2006) and certain classes of mathematical programs can be expressed, even in the presence of logical connectives. In terms of a solver strategy, they use binary search in tandem with lower and upper bounds to find the maximum (minimum). This is not unlike DPLL traces in knowledge compilation, which makes that technology the most accessible for propagating arbitrary semiring labels.
- Compositional settings are, of course, more intricate. Along with OMT, and classical iterative methods like *expectation maximization* (Koller and Friedman 2009), there are a number of recent approaches employing branch-and-bound search strategies to navigate between local and global consistency (Friesen and Domingos 2015). Which of these can be made amenable to compositions of SP programs remains to be seen however.

Overall, we believe the most promising first step is to limit the vocabulary of the logical language to propositions and constants (*i.e.*, 0-ary functions), which make appealing to knowledge compilation and OMT technology straightforward. It will also help us better characterize the complexity of the problems that SP attempts to solve.[7] At first glance, SP is seen to naturally capture #P-complete problems in the factorized setting, both in the finite case (Gomes, Sabharwal, and Selman 2009) and the infinite one (Dyer and Frieze 1988). In the non-factorized setting, many results from OMT and mathematical programming are inherited depending on the nature of the objective function and the domains of the program variables (Kannan and Monma 1978; Barrett et al. 2009; Sebastiani and Tomasi 2015). By restricting the language as suggested, the applicability of these results can be explored more thoroughly.

[6] Approaches like (Berre, Lonca, and Marquis 2014; Fargier, Marquis, and Schmidt 2013) on knowledge compilation for optimization and (Sanner and McAllester 2005) on circuits for linear constraints are compelling, but their applicability to non-trivial arithmetic and optimization problems remains to be explored.

6 Related Work

Semiring programming is related to efforts from different disciplines within AI, and we discuss representative camps.

Statistical modeling Formal languages for generative stochastic processes, such as Church (Goodman et al. 2008) and BLOG (Milch et al. 2005), have received a lot of attention in the learning community. Such languages provide mechanisms to compactly specify complex probability distributions, and appeal to sampling for inference.

Closely related to such proposals are probabilistic logic programming languages such as Problog (De Raedt, Kimmig, and Toivonen 2007) that extends Prolog with probabilistic choices and uses WMC for inference (Fierens et al. 2011). In particular, a semiring generalization of Problog, called aProbLog (Kimmig, Van den Broeck, and De Raedt 2011), was the starting point for our work and employs a semiring variation of WMC for inference (Kimmig, Van den Broeck, and De Raedt 2012), but none addresses matrix factorization, convex optimization and compositionality.

The use of semirings in machine learning is not new to aProbLog, see *e.g.*, (Goodman 1999), and programming languages such as Dyna (Eisner and Filardo 2011). Dyna is based on Datalog; our logical setting is strictly more expressive than Datalog and its extensions (*e.g.*, non-Horn fragment, constraints over reals). Dyna also labels proofs but not interpretations, as would SP (thus capturing weighted model counting, for example).

Constraints The constraints literature boasts a variety of modeling languages, such as Essence (Frisch et al. 2008), among others (Marriott et al. 2008; Van Hentenryck 2002). (See (Fontaine, Michel, and Van Hentenryck 2013), for example, for a proposal on combining heterogeneous solvers.) On the one hand, SP is more expressive from a logical viewpoint as constraints can be described using arbitrary formulas from predicate logic, and we address many problems beyond constraints, such as probabilistic reasoning. On the other hand, such constraint languages make it easier for non-experts to specify problems while SP, in its current form, assumes a background in logic. Such languages, then, would be of interest for extending SP's modeling features.

[7] Richer fragments have to be considered carefully to avoid undecidable properties (Boerger, Grädel, and Gurevich 1997).

A notable line of CSP research is by Bistarelli and his colleagues (Bistarelli et al. 1999; Bistarelli 2004). Here, semirings are used for diverse CSP specifications, which has also been realized in a CLP framework (Bistarelli, Montanari, and Rossi 2001). In particular, our account of compositionality is influenced by (Bistarelli 2004). Under some representational assumptions, SP and such accounts are related, but as noted, SP can formulate problems such as probabilistic inference in hybrid domains that does not have an obvious analogue in these accounts.

Optimization Closely related to the constraints literature are the techniques embodied in mathematical programming more generally. There are three major traditions in this literature that are related to SP. Modeling languages such as AMPL (Fourer, Gay, and Kernighan 1993) are fairly close to constraint modeling languages, and even allow parametrized constraints, which are ground at the time of search. The field of disciplined programming (Grant and Boyd 2008) supports features such as object-oriented constraints. Finally, relational mathematical programming (Apsel, Kersting, and Mladenov 2014) attempts to exploit symmetries in parametrized constraints.

From a solver construction perspective, these languages present interesting possibilities. From a framework point of view, however, there is little support for logical reasoning in a general way.

Knowledge representation Declarative problem solving is a focus of many proposals, including ASP (Brewka, Eiter, and Truszczyński 2011), model expansion (Mitchell and Ternovska 2005; Ternovska and Mitchell 2009), among others (Cadoli and Mancini 2007). These proposals are (mostly) for problems in NP, and so do not capture #P-hard problems like model counting and WMC. Indeed, the most glaring difference is the absence of weight functions over possible worlds, which is central to the formulation of statistical models. Weighted extensions of these formalisms, *e.g.*, (Baral, Gelfond, and Rushton 2009; Liu, Janhunen, and Niemelä 2012), are thus closer in spirit.

The generality of SP also allows us to instantiate many such proposals, including formalisms using linear arithmetic fragments (Belle, Passerini, and Van den Broeck 2015; Sebastiani and Tomasi 2015). Consider OMT for example. OMT can be used to express quantifier-free linear arithmetic sentences with a linear cost function, and a first-order structure that minimizes the cost function is sought. From a specification point of view, SP does not limit the logical language, does not require that objective functions be linear, and a variety of model comparisons, including counting, are possible via semirings. Compositionality in SP, moreover, goes quite beyond this technology.

Finally, there is a longstanding interest in combining different (logical) environments in a single logical framework, as seen, for example, in modular and multi-context systems (Lierler and Truszczynski 2015; Ensan and Ternovska 2015). In such frameworks, it would be possible to get a ILP program and ASP program to communicate their solutions, often by sharing atoms. In our view, Section 4 and these frameworks emphasize different aspects of compositionality. The SP scheme assumes the modeler will formalize a convex optimization problem and a SAT problem in the same programming language since they presumably arise in a single application (*e.g.*, a task and motion planner); this allows model reuse and enables transparency. In contrast, modular systems essentially treat diverse environments as black-boxes, which is perhaps easier to realize. On the one hand, it would be interesting to see whether modular systems can address problems such as combined inference and learning. On the other, some applications may require that different environments share atoms, for which our account on compositionality could be extended by borrowing ideas from modular systems.

7 Conclusions

In a nutshell, SP is a framework to declaratively specify three major concerns in AI applications:

- logical reasoning;
- discrete and continuous probabilistic inference;
- discrete and continuous optimization.

To its strengths, we find that SP is *universal* (in the above sense) and *generic* (in terms of allowing instantiations to particular logical languages and semirings). Thus, we believe SP represents a simple, uniform, modular and transparent approach to the model building process of complex AI applications.

SP comes with a rigorous semantics to give meaning to its programs. In that sense, we imagine future developments of SP would follow constraint programming languages and probabilistic programming languages in providing more intricate modeling features which, in the end, resort to the proposed semantics in the paper.

Perhaps the most significant aspect of SP is that it also allows us to go beyond existing paradigms as the richness of the framework admits novel formulations that combine theories from these different fields, as illustrated by means of a combined regression and probabilistic inference example. In the long term, we hope SP will contribute to the bridge between learning and reasoning.

References

[Andrei 2013] Andrei, N. 2013. *Nonlinear Optimization Applications Using the GAMS Technology*. Springer.

[Apsel, Kersting, and Mladenov 2014] Apsel, U.; Kersting, K.; and Mladenov, M. 2014. Lifting relational map-lps using cluster signatures. In *AAAI*, 2403–2409.

[Baral, Gelfond, and Rushton 2009] Baral, C.; Gelfond, M.; and Rushton, J. N. 2009. Probabilistic reasoning with answer sets. *TPLP* 9(1):57–144.

[Barrett et al. 2009] Barrett, C.; Sebastiani, R.; Seshia, S. A.; and Tinelli, C. 2009. Satisfiability modulo theories. In *Handbook of Satisfiability*. IOS Press. chapter 26, 825–885.

[Barrett, Fontaine, and Tinelli 2015] Barrett, C.; Fontaine, P.; and Tinelli, C. 2015. The SMT-LIB Standard: Version 2.5. Technical report, Department of Computer Science, The University of Iowa. Available at `www.SMT-LIB.org`.

[Belle, Passerini, and Van den Broeck 2015] Belle, V.; Passerini, A.; and Van den Broeck, G. 2015. Probabilistic inference in hybrid domains by weighted model integration. In *IJCAI*.

[Belle, Passerini, and Van den Broeck 2016] Belle, V.; Passerini, A.; and Van den Broeck, G. 2016. Towards component caching in hybrid domains. In *AAAI*.

[Berre, Lonca, and Marquis 2014] Berre, D. L.; Lonca, E.; and Marquis, P. 2014. On the complexity of optimization problems based on compiled NNF representations. *CoRR* abs/1410.6690.

[Bistarelli et al. 1999] Bistarelli, S.; Montanari, U.; Rossi, F.; Schiex, T.; Verfaillie, G.; and Fargier, H. 1999. Semiring-based CSPs and valued CSPs: Frameworks, properties, and comparison. *Constraints* 4(3):199–240.

[Bistarelli, Montanari, and Rossi 2001] Bistarelli, S.; Montanari, U.; and Rossi, F. 2001. Semiring-based constraint logic programming: syntax and semantics. *TOPLAS* 23(1):1–29.

[Bistarelli 2004] Bistarelli, S. 2004. *Semirings for Soft Constraint Solving and Programming*. SpringerVerlag.

[Boerger, Grädel, and Gurevich 1997] Boerger, E.; Grädel, E.; and Gurevich, Y. 1997. *The classical decision problem*. Springer Verlag.

[Brewka, Eiter, and Truszczyński 2011] Brewka, G.; Eiter, T.; and Truszczyński, M. 2011. Answer set programming at a glance. *Communications of the ACM* 54(12):92–103.

[Cadoli and Mancini 2007] Cadoli, M., and Mancini, T. 2007. Combining relational algebra, sql, constraint modelling, and local search. *Theory and Practice of Logic Programming* 7:37–65.

[Chavira and Darwiche 2008] Chavira, M., and Darwiche, A. 2008. On probabilistic inference by weighted model counting. *Artif. Intell.* 172(6-7):772–799.

[Chistikov, Dimitrova, and Majumdar 2015] Chistikov, D.; Dimitrova, R.; and Majumdar, R. 2015. Approximate counting in smt and value estimation for probabilistic programs. In *TACAS*, volume 9035. 320–334.

[De Raedt, Kimmig, and Toivonen 2007] De Raedt, L.; Kimmig, A.; and Toivonen, H. 2007. Problog: A probabilistic prolog and its application in link discovery. In *Proc. IJCAI*, 2462–2467.

[Ding, Li, and Jordan 2010] Ding, C. H. Q.; Li, T.; and Jordan, M. I. 2010. Convex and semi-nonnegative matrix factorizations. *IEEE Trans. Pattern Anal. Mach. Intell.* 32(1):45–55.

[Dyer and Frieze 1988] Dyer, M. E., and Frieze, A. M. 1988. On the complexity of computing the volume of a polyhedron. *SIAM J. Comput.* 17(5):967–974.

[Eisner and Filardo 2011] Eisner, J., and Filardo, N. W. 2011. Dyna: Extending Datalog for modern AI. In *Datalog Reloaded*, LNCS. Springer. 181–220.

[Enderton 1972] Enderton, H. 1972. *A mathematical introduction to logic*. Academic press New York.

[Ensan and Ternovska 2015] Ensan, A., and Ternovska, E. 2015. Modular systems with preferences. In *IJCAI*, 2940–2947.

[Fargier, Marquis, and Niveau 2013] Fargier, H.; Marquis, P.; and Niveau, A. 2013. Towards a knowledge compilation map for heterogeneous representation languages. In *IJCAI*, 877–883.

[Fargier, Marquis, and Schmidt 2013] Fargier, H.; Marquis, P.; and Schmidt, N. 2013. Semiring labelled decision diagrams, revisited: Canonicity and spatial efficiency issues. In *IJCAI*.

[Fierens et al. 2011] Fierens, D.; Van den Broeck, G.; Thon, I.; Gutmann, B.; and De Raedt, L. 2011. Inference in probabilistic logic programs using weighted CNF's. In *UAI*, 211–220.

[Fontaine, Michel, and Van Hentenryck 2013] Fontaine, D.; Michel, L.; and Van Hentenryck, P. 2013. Model combinators for hybrid optimization. In *CP*, 299–314.

[Fourer, Gay, and Kernighan 1993] Fourer, R.; Gay, D. M.; and Kernighan, B. W. 1993. *AMPL : a modeling language for mathematical programming*. South San Francisco: Scientific Press.

[Freuder and Mackworth 2006] Freuder, E. C., and Mackworth, A. K. 2006. Constraint satisfaction: An emerging paradigm. *Handbook of Constraint Programming* 13–28.

[Freuder 1997] Freuder, E. 1997. In pursuit of the holy grail. *Constraints* 2(1):57–61.

[Friesen and Domingos 2015] Friesen, A., and Domingos, P. 2015. Recursive decomposition for nonconvex optimization. In *IJCAI*.

[Frisch et al. 2008] Frisch, A. M.; Harvey, W.; Jefferson, C.; Hernández, B. M.; and Miguel, I. 2008. Essence : A constraint language for specifying combinatorial problems. *Constraints* 13(3):268–306.

[Gomes, Sabharwal, and Selman 2009] Gomes, C. P.; Sabharwal, A.; and Selman, B. 2009. Model counting. In *Handbook of Satisfiability*. IOS Press.

[Goodman et al. 2008] Goodman, N. D.; Mansinghka, V. K.; Roy, D. M.; Bonawitz, K.; and Tenenbaum, J. B. 2008. Church: a language for generative models. In *Proc. UAI*, 220–229.

[Goodman 1999] Goodman, J. 1999. Semiring parsing. *Computational Linguistics* 25(4):573–605.

[Grant and Boyd 2008] Grant, M., and Boyd, S. 2008. Graph implementations for nonsmooth convex programs. In *Recent Advances in Learning and Control*, Lecture Notes in Control and Information Sciences. Springer-Verlag Limited. 95–110.

[Halmos 1950] Halmos, P. 1950. Measure theory. *Van Nostrad Reinhold Company*.

[Halpern 1990] Halpern, J. 1990. An analysis of first-order logics of probability. *Artificial Intelligence* 46(3):311–350.

[Hindi 2004] Hindi, H. 2004. A tutorial on convex optimization. In *American Control Conference*, volume 4.

[Hooker and Lama 1999] Hooker, J. N., and Lama, M. A. O. 1999. Mixed logical-linear programming. *Discrete Applied Mathematics* 96-97:395–442.

[Kannan and Monma 1978] Kannan, R., and Monma, C. L. 1978. *On the computational complexity of integer programming problems*. Springer.

[Kautz 2004] Kautz, H. 2004. Toward a universal inference engine. In *LPNMR*, volume 2923 of *LNCS*. Springer Berlin Heidelberg. 2–2.

[Kimmig, Van den Broeck, and De Raedt 2011] Kimmig, A.; Van den Broeck, G.; and De Raedt, L. 2011. An algebraic prolog for reasoning about possible worlds. In *Proc. AAAI*.

[Kimmig, Van den Broeck, and De Raedt 2012] Kimmig, A.; Van den Broeck, G.; and De Raedt, L. 2012. Algebraic model counting. *CoRR* abs/1211.4475.

[Koller and Friedman 2009] Koller, D., and Friedman, N. 2009. *Probabilistic Graphical Models - Principles and Techniques*. MIT Press.

[Kuich 1997] Kuich, W. 1997. Semirings and formal power series: their relevance to formal languages and automata. In *Handbook of formal languages, Vol. 1*. Berlin: Springer. 609–677.

[Lierler and Truszczynski 2015] Lierler, Y., and Truszczynski, M. 2015. An abstract view on modularity in knowledge representation. In *AAAI*, 1532–1538.

[Liu, Janhunen, and Niemelä 2012] Liu, G.; Janhunen, T.; and Niemelä, I. 2012. Answer set programming via mixed integer programming. In *KR*.

[Marriott et al. 2008] Marriott, K.; Nethercote, N.; Rafeh, R.; Stuckey, P.; Garcia de la Banda, M.; and Wallace, M. 2008. The design of the zinc modelling language. *Constraints* 13(3):229–267.

[Milch et al. 2005] Milch, B.; Marthi, B.; Russell, S. J.; Sontag, D.; Ong, D. L.; and Kolobov, A. 2005. BLOG: Probabilistic models with unknown objects. In *Proc. IJCAI*, 1352–1359.

[Mitchell and Ternovska 2005] Mitchell, D. G., and Ternovska, E. 2005. A framework for representing and solving NP search problems. In *AAAI*, 430–435.

[Nieuwenhuis and Oliveras 2006] Nieuwenhuis, R., and Oliveras, A. 2006. On sat modulo theories and optimization problems. In *Theory and Applications of Satisfiability Testing-SAT 2006*. Springer. 156–169.

[Richardson and Domingos 2006] Richardson, M., and Domingos, P. 2006. Markov logic networks. *Machine learning* 62(1):107–136.

[Sankaranarayanan, Chakarov, and Gulwani 2013] Sankaranarayanan, S.; Chakarov, A.; and Gulwani, S. 2013. Static analysis for probabilistic programs: inferring whole program properties from finitely many paths. *ACM SIGPLAN Notices* 48(6).

[Sanner and McAllester 2005] Sanner, S., and McAllester, D. A. 2005. Affine algebraic decision diagrams (aadds) and their application to structured probabilistic inference. In *IJCAI*, 1384–1390.

[Sebastiani and Tomasi 2015] Sebastiani, R., and Tomasi, S. 2015. Optimization modulo theories with linear rational costs. *ACM Trans. Comput. Logic* 16(2):12:1–12:43.

[Srivastava et al. 2014] Srivastava, S.; Fang, E.; Riano, L.; Chitnis, R.; Russell, S. J.; and Abbeel, P. 2014. Combined task and motion planning through an extensible planner-independent interface layer. In *ICRA*, 639–646.

[Ternovska and Mitchell 2009] Ternovska, E., and Mitchell, D. G. 2009. Declarative programming of search problems with built-in arithmetic. In *Proc. IJCAI*, 942–947.

[Thrun, Burgard, and Fox 2005] Thrun, S.; Burgard, W.; and Fox, D. 2005. *Probabilistic Robotics*. MIT Press.

[Van Hentenryck 2002] Van Hentenryck, P. 2002. Constraint and integer programming in OPL. *INFORMS Journal on Computing* 14(4):345–372.